GPCD Weather Normalization Methodology

Policy

Planning

Research

Evaluation

By

Anil Bamezai, Ph.D.
310.314.7691
bzi@mindspring.com

December 13, 2011

WESTERN
POLICY
RESEARCH

Final Report Submitted to the California Urban Water Conservation Council

171 Pier Avenue
Suite 256
Santa Monica
California 90405

Preface

This paper develops a methodology for weather normalizing agency-level production data, required for verifying compliance with the Memorandum of Understanding under the GPCD Compliance Option. It may also be useful for testing compliance with the SBx7-7 legislation, although this new legislation has many more components which will be more fully developed in the future by the California Department of Water Resources.

Contents

Figures and Tables

Executive Summary

Council members that opt to reach their water conservation goals using the GPCD Compliance Option must test compliance by tracking annual consumption in terms of gallons per capita per day (GPCD). Over time, a supplier's GPCD may change because of a combination of factors including, conservation, water rates, technology, economic conditions, customer tastes, and weather. The GPCD Compliance Option only allows for adjustments to GPCD due to unusual weather since weather is beyond a water supplier's control. Most of the other factors are partially within a supplier's sphere of influence and in fact represent tools that can be used to reduce GPCD over time. Thus, no need arises in principle for normalizing GPCD on account of these non-weather related factors.

This study was undertaken to develop a weather-normalization methodology, and in the process address several questions including: (1) what data source should suppliers use for obtaining reliable weather information; (2) how should one account for weather's impact upon water demand, and the variation in this relationship across suppliers; and (3) should the preferred metrics for capturing weather be reference ETo and rainfall or temperature and rainfall? Data from a total of 18 diverse suppliers were assembled in two phases to address these questions.

Regarding the first question, there was a great deal of interest in assessing whether PRISM can serve as a one-stop shop for weather data? The analyses presented here suggest a clear yes to this question. PRISM is a powerful new tool developed by the DWR that can provide temperature, rainfall, and reference ETo data from 1990 into the future for all regions of California. Extensive comparisons between PRISM data and data from other sources, such as NOAA and CIMIS, found them to be highly correlated. For example, temperature and rainfall from NOAA and PRISM show a correlation of 0.96 and 0.94, respectively, across the suppliers included in this study (Section 3).

The second key question that this study had to deal with was accounting for variation in weather response across different suppliers. From the outset we recognized that each supplier's mix of weather-sensitive and weather-insensitive end uses is different; therefore, the impact of weather on total production cannot be identical across suppliers. Suppliers, however, do not always have good data to isolate these two types of end uses, so a weather normalization scheme that relies on the availability of such disaggregate data is likely to run into difficulty. Alternatively, one could collect supplier characteristics that correlate with weather-sensitive end-uses, such as irrigated landscape per capita, intensity of commercial air-conditioning, and so on, but these too are difficult to obtain in practice. Thus, we focused on a supplier's peaking factor as a way of scoring how suppliers rank relative to one another in terms of the proportion of total use that is accounted for by weather-sensitive end uses. For peaking factor to work as a scoring variable, however, it must have a consistent relationship with the variation in weather impacts across suppliers.

To test whether such a relationship exists we first made sure that our sample of 18 suppliers exhibited a range of peaking factors that matches California's as a whole. California has water suppliers with summer-to-winter monthly production ratios (peaking factor) extending to 5 and a few beyond 5. Our sample includes suppliers with peaking factors that range from a minimum of 1.6 to a maximum of 6, from different locations in the state (coastal, inland, north, south, central).

Section 4 develops a theoretical model which suggests that weather's impact should scale nonlinearly with respect to a supplier's peaking factor. Expectations from this theoretical model were extensively tested and found to corroborate well with the data. Taken together, all of these empirical findings suggest that using a supplier's peaking factor to implicitly score a supplier's mix of weather sensitive and insensitive uses is a viable modeling strategy—and therefore a viable weather normalization strategy. Once this insight was tested and established it was relatively straightforward to pool data across all 18 suppliers and estimate peaking-factor dependent weather impacts with a high level of statistical precision.

Given the key role that peaking factor plays in our methodology, for which working with monthly production data is necessary, it should be self-evident that weather normalization at the annual, instead of monthly level, is unlikely to succeed. There was interest in this question—can weather normalization work at the annual level—since it would considerably simplify the data collection requirements and the computations involved. Our analyses indicate that this is not a feasible option.

Finally, we performed several sensitivity analyses to determine whether reference ETo and rainfall or temperature and rainfall should be the preferred metrics for capturing weather. While both approaches yield comparable results, we feel that weather normalization based upon rainfall-adjusted reference ETo is likely to yield more reliable results than one based upon temperature and rainfall. While testing the relationship between supplier-specific weather impacts and peaking factors, we found it to be tightest when weather impacts were captured using rainfall-adjusted reference ETo, less so when these impacts were captured separately using temperature and rainfall. Thus a methodology that does not require estimating the independent effect of rainfall on production is likely to be more reliable, and the weight of the empirical evidence presented in Section 4 certainly supports this assertion.

Acknowledgments

We are very grateful to all the water suppliers' representatives that volunteered their agency's data for this study, helped us identify suitable sources of weather data, and helped us work through data issues as and when they arose. Their diligence was a key precursor to this project's success.

1. Introduction

Council members that opt to reach their water conservation goals using the GPCD Compliance Option must test compliance by tracking annual consumption in terms of gallons per capita per day (GPCD). Under this option, suppliers need to report their GPCD every two years with the overall goal of reducing their 2018 GPCD by 18% relative to their baseline GPCD.

Over time, a supplier's GPCD may change because of a combination of factors including, conservation, water rates, technology, economic conditions, customer tastes, and weather. The GPCD Compliance Option only allows for adjustments to GPCD due to unusual weather since weather is beyond a water supplier's control. It is easy to see that weather may be too hot or too cold during a compliance year relative to the baseline years, which must be taken into account while assessing compliance. Most of the other factors are partially within a supplier's sphere of influence and in fact represent tools that can be used to reduce GPCD over time. Thus, no need arises in principle for normalizing GPCD on account of these non-weather related factors. The question about unusual economic conditions, however, remains worrisome. The current economic environment has depressed water demand in most areas. If these conditions prevail for an extended period it may lull suppliers into believing they are easily meeting their interim targets, only to be unpleasantly surprised as water demands rebound later in the decade.

That said this study only focuses on how to normalize GPCD for unusual weather. The question about how to deal with unusual economic conditions is beyond this study's scope. Note, that while we commonly speak of weather normalizing GPCD, weather only influences water demand, not population. Thus, throughout this report, when we speak of weather normalizing GPCD, what we mean is weather normalizing production first—production is used as a proxy for demand—then expressing these normalized production data in GPCD terms.

1.1 Key Study Questions

This study was undertaken to develop a weather-normalization methodology, and in the process address several questions including:

1. Does the source of weather data matter?
2. Is weather normalization of aggregate annual production data feasible, or are disaggregate monthly data necessary?
3. Is there a performance difference between using reference ETo and rainfall over temperature and rainfall to perform the weather normalization?

4. To what extent does water demand's response to weather vary across suppliers?

This study proceeded in two phases. In phase I, the California Urban Water Conservation Council (CUWCC) provided production data going back in time for 10 California water suppliers to develop and test an appropriate weather normalization methodology. During the initial analyses, the composition of this initial sample was discovered to be somewhat skewed toward suppliers with the summer-to-winter monthly production ratios (peaking factor) under 3. California includes suppliers with monthly peaking factors up to 5, and some beyond 5.

Peaking factor is of particular interest to us as a modeling variable since we do not have detailed information about supplier characteristics that correlate with weather-sensitive end uses. In the absence of detailed end-use characteristics, a supplier's peaking factor becomes an easy-to-observe proxy for the proportion of total demand that is accounted for by irrigation and other weather-sensitive end-uses, such as industrial chillers. By characterizing suppliers according to their respective peaking factors, a general weather normalization methodology can in principle be developed that works for different types of suppliers (for example, those located on the coast versus inland, those with greater or lesser landscape per capita, or those with greater or lesser CII use per capita). Do the data support use of peaking factors as a modeling input? This is a key question, addressed later. We are, however, not advocating the use of peaking factors as a proxy for water use efficiency. Peaking factors contain useful information about the weather-sensitive portion of total demand, but they are weakly correlated with the level of water use efficiency achieved by a supplier.

The phase I sample was used extensively to principally address the first question, and to a lesser extent the second and third questions. This initial sample was then supplemented in Phase II with 9 additional suppliers with peaking factors extending to 6, to address more fully questions two through four.

For the purpose of this study, we have suppressed the identity of the suppliers that kindly volunteered their data. Since these suppliers may improve their production and population data to make them consistent with Department of Water Resources' (DWR) data preparation guidelines[1], and also may choose a different set of years to define their baseline, their officially reported baseline GPCD estimates may differ somewhat from ours. This is not a problem. But, we thought it best to avoid confusion that competing published estimates may generate.

[1] See "Methodologies for Calculating Baseline and Compliance Urban Per Capita Water Use," issued by the California Department of Water Resources.

1.2 Establishing Baseline GPCD

To weather normalize consumption in the compliance years, one must have a benchmark to normalize to—in other words, one must be able to define a given area's normal or average weather. According to the GPCD Compliance Option, water retailers need to select 10 consecutive years of production and population data (1997 through 2006) to establish their baseline; exceptions are permitted for suppliers that signed the MOU prior to 1997 as outlined in the MOU Compliance Policy.

For the purpose of compliance then, the benchmark becomes the "average" weather that prevailed during the years that enter into the determination of the baseline GPCD. Some suppliers may have weather data going back several decades, but not all these years should be used to establish "average" weather, only the years that enter the baseline GPCD calculation.

1.3 Council's GPCD Compliance Option Versus SBx7-7

The Council's GPCD Compliance Option has many similarities, but also some differences, with the SBx7-7 legislation's compliance approach. The latter is covered in detail in DWR's Technical Methodologies document cited earlier. Key differences are as follows:

- SBx7-7 allows suppliers to choose between one of four methods to select their final (2020) targets. The Council allows only one method, an 18% reduction in GPCD by 2018 relative to the baseline.
- SBx7-7 requires compliance testing in the interim year 2015 and final year 2020, while the Council requires GPCD reporting every other year.
- SBx7-7 offers a bit more flexibility with respect to time period and number of years to be included in the baseline.
- SBx7-7 permits weather normalization of compliance year GPCD, just as the Council does, but the former additionally permits adjustment to compliance year GPCD for unusual economic conditions.
- SBx7-7 permits suppliers to form a regional group for the purpose of compliance testing, while the Council's GPCD Compliance Option is available only to individual suppliers.

1.4 Sources of Weather Data

Reliable weather data remains a key factor in any GPCD normalization methodology (in addition to having reliable production and population data). At the beginning of this project, we assumed that good quality weather data would be available for all the test agencies, but in fact this proved to be a challenge. It is worth dwelling a bit on this subject.

It is well known that saturation of CIMIS stations in urban areas is patchy. To overcome this shortcoming, DWR has developed a tool called SpatialCIMIS, which presumably can provide model-interpolated reference ETo information for any part of California. We had expected to use data from this tool to fill gaps in the coverage of the CIMIS station network. But, unfortunately, we learned that this tool only goes back to 2003, insufficient to match the 10 year baseline that extends back into the mid-1990s.

DWR also has another tool called PRISM that provides historical model interpolated weather data including temperature, rainfall, and reference ETo for each 4x4 kilometer tile in the State of California. There are slightly over 26,000 such tiles. Reference ETo data available from PRISM, however, are based on the Hargreaves-Samani (HS) model that uses only temperature and latitude to predict reference ETo. On the other hand, reference ETo available from CIMIS is based upon the Penman-Montieth (PM) model that includes solar radiation, wind speed and relative humidity in addition to temperature. Because the HS model only uses temperature and latitude to generate reference ETo, the PRISM tool developers compared both the HS and PM model estimates for tiles where both were available to determine correction factors that were then used to convert HS ETo estimates into PM-equivalent ETo estimates.

The PRISM data at present are available from 1990 through 2010. DWR is planning to extend the PRISM data to present time and also to develop a user interface that would allow any user in the state to download weather data appropriate to their service area.

To assess whether PRISM can serve as a one-stop shop for weather data, we compared the efficacy of these data with weather data obtained from other sources as well. A key source of temperature and rainfall data is the National Oceanic and Atmospheric Administration's (NOAA) data archives available through the Western Regional Climate Center (http://www.wrcc.dri.edu/). While the coverage of these stations appears to be very impressive on a map, causing one at first blush to assume that a suitable station could probably be found for just about every California water supplier, many stations are plagued by missing or incomplete data. So obtaining reliable weather data from this source is challenging, which is why a one-stop source, such as PRISM remains so attractive. But, how good are the PRISM data?

To assess how sensitive weather normalization is with respect to the source of weather data, we assembled temperature and rainfall data for all the 10 Phase I suppliers using NOAA stations, sometimes supplemented with data from close-by CIMIS stations to fill in bad or missing NOAA data. The PRISM based temperature, rainfall, and reference ETo data are of course available for all Phase I and II suppliers. Thus, first we use Phase I suppliers to assess how weather data correlate across alternative data sources. Then the full sample is employed, using PRISM weather data, to assess water demand's relationship with weather, and to assess how this relationship varies across suppliers with different peaking factors.

2. Analytic Approach

2.1 Key Steps

The purpose of weather normalizing compliance-year GPCD is to ask the following question—what would GPCD *have been* during the compliance year had weather been the same as it was on average during the baseline period? Answering this question required us, and will require the Council to work through the following six steps for those suppliers that choose the GPCD Compliance Option:

1. Select a baseline period. Suppliers would have to furnish monthly (potable) water production data for their baseline period (1997 through 2006 per the MOU).
2. Estimate average weather for the baseline period. Suppliers would have to select and inform the Council about which PRISM tiles they wish to use for capturing weather in their service area.
3. Use models developed in this study to estimate how production responds to deviations in weather. The Council would complete this step. No additional data from suppliers would be required.
4. Estimate weather deviations during the compliance year relative to the baseline. The Council would complete this step based upon PRISM tiles selected in Step 2.
5. Combine information from steps 3 and 4 to derive weather normalized production during the compliance year. The Council would complete this step.
6. Divide compliance-year production derived in step 5 by population and days in the year to derive weather-normalized GPCD. The Council would complete this step.

The main goal of this study is Step 3—estimating the strength of the relationship between production and weather, which can then be used to derive weather-normalized production from actual production[2].

As discussed in detail later, the annual model was not successful in detecting a statistically significant relationship between production and weather. Therefore, weather normalization of annual production, while attractive for its simplicity, was determined to be infeasible.

2.2 Modeling Strategy

To reliably estimate how production reacts to changes in weather one must also account for other factors that influence production within a year, and over the course of many years. Key factors here include season

[2] The Council has also prepared guidelines for suppliers to follow that wish to undertake the above six-step approach, including Step 3, on their own.

and population which we are able to account for (the former only in the monthly models). The explanatory power of just these three factors (season, population, and weather deviation relative to the baseline) is so high that it is unlikely that important confounding factors remain unaccounted for. However, to test for this possibility we included time trend variables for each agency to capture non-cyclical time trends in production not captured by population. These trend variables added no significant explanatory power to the models.

The models are estimated using 10 years of data prior to 2005. It is not theoretically necessary to restrict these models to only 10 years. Obtaining statistically significant correction factors was not a problem using 10 years of data. Including years beyond what is necessary to obtain statistical significance increases the risk of introducing unaccounted for time-varying confounding factors, so we did not. In any event, including data past 2007 for model estimation was decidedly out of the question because of the economic recession. What we have instead done is use the years 2006, 2007 and 2008 to forecast the outcome of weather normalization and to conduct sensitivity analyses. We derive weather-normalized GPCDs for each agency in these years using the temperature and rainfall approach and contrast how the results change when normalization is attempted using reference ETo and rainfall.

For the statistically inclined, Appendix A provides greater detail about the model structure, results of the statistical estimation, and sensitivity analyses. Our analyses exclude Supplier #8 because their production and population data exhibit year-over-year anomalies that the supplier was unable to explain or fix. Study results are not sensitive to the exclusion of this supplier, however.

But first we have to establish the efficacy of PRISM weather data. This was done by working with data collected from phase I suppliers, for whom we assembled weather data from both NOAA/CIMIS stations as well as from PRISM.

3 Assessing Efficacy of PRISM Data

3.1 Phase I Test Suppliers

Table 1 shows the data we received from each of the Phase I suppliers, and the steps we had to undertake to create a temperature and rainfall data series from sources other than PRISM for the purpose of comparison.

In most cases, we used the period 1995 through 2004 to set the baseline. For developing and testing a weather normalization methodology, this is acceptable even though suppliers may make alternative choices for official reporting to the CUWCC or DWR. A couple of suppliers did not give us data going back to 1995. For these suppliers, baseline GPCDs could not be estimated using 10 years of production data, which may affect the accuracy of their baseline estimates, but otherwise posed no modeling difficulty.

With respect to weather data we have tapped three sources, NOAA, CIMIS, and PRISM. Since PRISM offers full coverage of the state, we have temperature, rainfall, and reference ETo from this source for all sampled suppliers, which is the main reason why this data source is so appealing. From the NOAA archives, we were able to piece together an independent temperature and rainfall data series for 9 suppliers. Bad and missing weather data in the primary NOAA station were imputed using data from other close-by NOAA or CIMIS station(s). In the case of one supplier, however, we were unable to obtain any NOAA data at all so we relied only on CIMIS to obtain temperature and rainfall data. CIMIS and NOAA temperature and rainfall data are highly correlated.

Appropriate CIMIS stations could be identified only in the case of 4 out of the 10 phase I suppliers, proof of CIMIS's patchy urban coverage and unsuitability as a source of consistent weather data on a statewide basis.

The good news is that monthly temperature and rainfall data from NOAA and PRISM appear to be highly correlated—0.96 and 0.94 respectively. This strengthens the *a priori* case for relying on PRISM data for the weather normalization of production data.

Table 2 shows the agency-by-agency correlation between monthly NOAA and PRISM weather data, which are all very good.

Table 1 Phase I data structure and quality

Supplier ID	Type of Data	Baseline Period	CIMIS data available?	NOAA data available?	PRISM data available?	Comments
1	Monthly	1998-2004	No	Yes	Yes	Agency did not provide data going back to 1995. Therefore, only 7 years enter baseline GPCD determination. Missing weather data in primary NOAA station imputed using secondary station.
2	Monthly	1995-2004	Yes	No	Yes	ALL appropriate NOAA stations had unsalvageable data problems. CIMIS station used instead.
3	Monthly	1995-2004	Yes	Yes	Yes	Good correlation between CIMIS and NOAA temperature and rainfall data.
4	Monthly	1995-2004	Yes	Yes	Yes	CIMIS data available with discontinuities, used mainly to impute missing NOAA station data.
5	Monthly	1995-2004	No	Yes	Yes	High quality NOAA data found and averaged over 3 stations spanning service area.
6	Monthly	1995-2004	No	Yes	Yes	Fairly high quality NOAA data found from single station.
7	Monthly	1995-2004	No	Yes	Yes	CIMIS data available with severe discontinuities, used mainly to impute missing NOAA station data.
8	Monthly	1999-2004	No	Yes	Yes	Agency did not provide data going back to 1995. Production and population data have anomalies—agency dropped from analyses.
9	Monthly	1995-2004	No	Yes	Yes	Missing weather data in primary NOAA station imputed using secondary station.
10	Monthly	1995-2004	Yes	Yes	Yes	Closest two CIMIS stations suffer from discontinuities, but another nearby CIMIS station, highly correlated with the closest two, available. NOAA station available with high quality data.

Table 2 Correlation between monthly NOAA and PRISM data

Supplier ID	Temperature correlation	Rainfall correlation
1	0.99	0.99
2	0.98	0.96
3	0.93	0.88
4	0.98	0.84
5	0.99	0.99
6	0.98	0.97
7	0.96	0.98
9	0.99	0.98
10	0.95	0.92
Overall	0.96	0.94

4 Estimating the Relationship between Production and Weather

4.1 The Importance of Peaking Factor

The previous section established the feasibility of using PRISM weather data to weather normalize monthly production. We now turn to describing the details of a methodology with statewide applicability for accomplishing this goal.

We recognize that each supplier's mix of weather-sensitive and weather-insensitive end uses is different; therefore, the impact of weather on total production cannot be identical across suppliers. Suppliers, however, do not always have good data to isolate these two types of end uses, so a weather normalization scheme that relies on the availability of such disaggregate data cannot succeed. Alternatively, one could collect supplier characteristics that correlate with weather-sensitive end-uses, such as irrigated landscape per capita, intensity of commercial air-conditioning, and so on, but these too are difficult to obtain in practice. Thus, we focused on a supplier's peaking factor as a way of scoring how suppliers rank relative to one another in terms of the proportion of total use that is accounted for by weather-sensitive end uses. For peaking factor to work as a scoring variable, however, it must have a consistent relationship with the variation in weather impacts across suppliers. Whether such a relationship exists was tested in several ways by first running separate models for each supplier.

As mentioned earlier, California has water suppliers with summer-to-winter monthly production ratios (peaking factor) extending to 5 and a few beyond 5, but our phase I sample only included suppliers with peaking factors under 3. While this was adequate for assessing efficacy of PRISM weather data, it was inadequate as far as developing a methodology with statewide applicability was considered. The strength of production's relationship with weather ought to vary as a function of the proportion of total production that is accounted for by weather sensitive end uses. To test and account for this possibility, and thereby arrive at a robust framework with statewide applicability, the original sample was supplemented with additional suppliers with significantly higher peaking factors. Thus, Section 4's analyses are based upon a total of 18 suppliers, 9 from phase I and an additional 9 from phase II.

While it is intuitive to expect that weather's impact upon production ought to scale positively with respect to a supplier's peaking factor, the exact nature of this relationship is not obvious. Should these weather effects scale linearly, logarithmically, or in some other way? To shed light on this, we offer a thought experiment.

Imagine several suppliers with 1 unit of non-weather sensitive demand, but progressively higher levels of weather sensitive demands (Table 3). Assume supplier B, at the peak of summer, uses an additional 1 unit of water in a normal year to meet its weather sensitive demand. If so, supplier B would be deemed to have a peaking factor of 2. Supplier C uses an additional 2 units of water in a normal year to meet its maximum summer demand, leading to a peaking factor of 3, and so on. Now let's ask by how much total summer demand would change if reference ETo were higher by 10% relative to normal. In the case of supplier B, weather sensitive use would increase from 1 to 1.1 units, leading to a 5% (2.1/2) increase in total summer demand relative to a normal year. In the case of supplier C, weather sensitive use would increase from 2 to 2.2 units, increasing total summer demand by 6.7% (3.2/3), and so on.

From this thought experiment, it is also easy to see that weather's impact is not a function of average GPCD holding peaking factor constant. For example, if another hypothetical supplier B[1] had an insensitive water demand of 2 units and summer weather sensitive demand of an additional 2 units, its average GPCD would be double that of supplier B, but the peaking factor would be the same. A 10% hotter year would raise their total summer demand by exactly the same 5% even though their average GPCDs were quite different. In other words, for scaling the impact of weather on production, the critical variable is peaking factor, not average GPCD.

Table 3 Theoretical relationship between peaking factor and weather's impact on total demand

Supplier	Weather insensitive demand	Summer weather sensitive demand	Total summer demand	Peaking factor	Percent increase in total summer demand	1-(1/peaking factor)
A	1 unit	0 unit	1 units	1	0.0%	0.00
B	1	1	2	2	5.0%	0.50
C	1	2	3	3	6.7%	0.67
D	1	3	4	4	7.5%	0.75
E	1	4	5	5	8.0%	0.80
F	1	5	6	6	8.3%	0.83
G	1	6	7	7	8.6%	0.86
H	1	7	8	8	8.8%	0.88
I	1	8	9	9	8.9%	0.89
J	1	9	10	10	9.0%	0.90

Figure 1 shows a plot of the increase in total summer demand by peaking factor that would result from weather being 10% hotter than normal. One can see that the relationship is highly nonlinear. At a peaking factor of 1 (indicating zero weather sensitive demand), the impact of weather is zero as one would expect. And as peaking factors continue to increase, indicating a higher proportion of weather-sensitive uses, the impact of a 10% increase in reference ETo plateaus at 10%. In other words, if total demand comprised only of irrigation and weather was 10% hotter than normal, then total demand would also be 10% greater than normal.

Modeling this nonlinear relationship between peaking factor and weather's impact on demand would be significantly easier if the peaking factor variable could somehow be transformed to linearize the relationship. This is accomplished by creating a new variable (1-(1/peaking factor)), which as Table 3 shows scales linearly with the expected percentage impacts on total production. This fundamental insight about how to transform peaking factors proved very useful for the purpose of model specification.

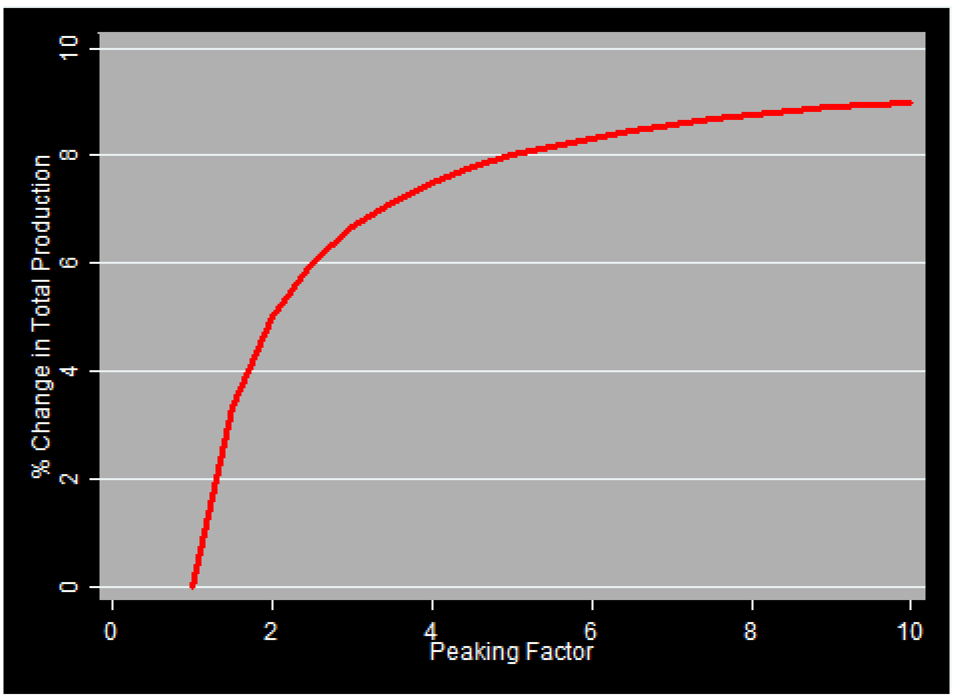

Figure 1 Nonlinear relationship between weather's impact and peaking factor

But, do the actual data support expectations from the above thought experiment? Yes, they do. To test for this, we ran separate models for each supplier two different ways; first to estimate how rainfall adjusted reference ETo impacts production; second, how temperature and rainfall impact production. The model coefficients (one per supplier) from the first set of models are plotted by the supplier's peaking factor (Figure 2) and by the transformed peaking factor (Figure 3). These supplier-specific models also yield the peaking factors. The model coefficients can be interpreted as the percent change in monthly production from a 1 inch deviation in rainfall adjusted reference ETo. At a coefficient equal to 0.1, a 1 inch positive deviation can be expected to raise monthly production by roughly 10 percent ($e^{0.1}$-1), and so on.

As expected, Figure 2 indicates a nonlinear relationship, which becomes linear in Figure 3 (with a correlation equal to 0.77). It can also be verified that a best-fit straight line estimated for Figure 3's data predicts a zero weather impact at a peaking factor of one, or at a transformed peaking factor of zero (Table 4 shows that the estimated intercept is statistically insignificant from zero). Our simple thought experiment anticipates all three outcomes. These simple plots and tests thus bolster our confidence in the use of transformed peaking factors to scale weather impacts across suppliers with different mixes of end uses and resultant demand profiles.

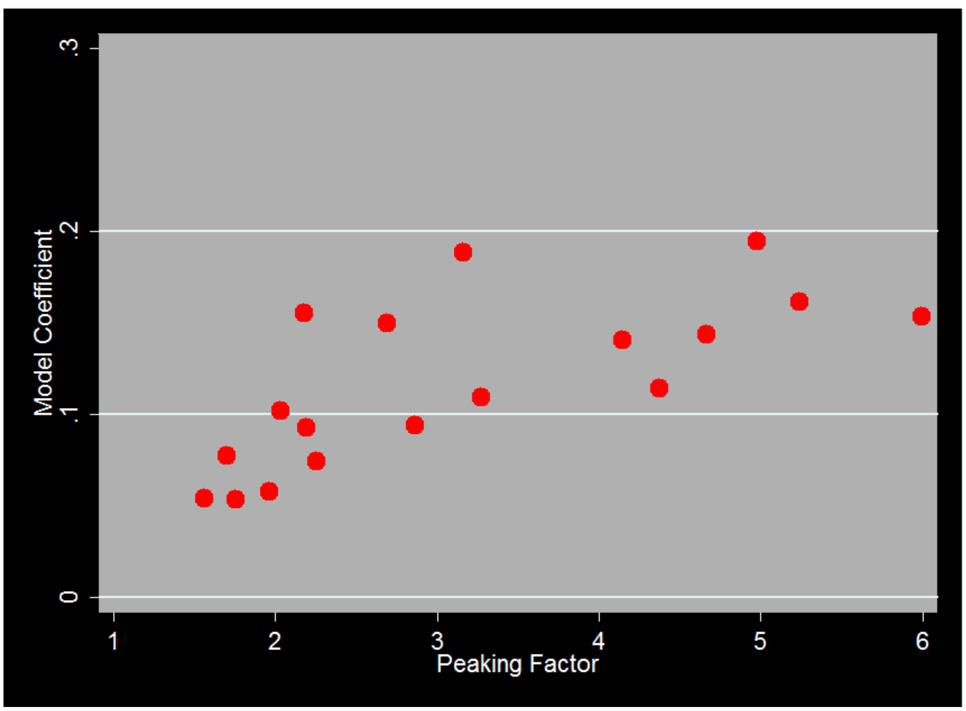

Figure 2 Supplier specific weather impacts as captured by rainfall adjusted reference ETo plotted by peaking factor

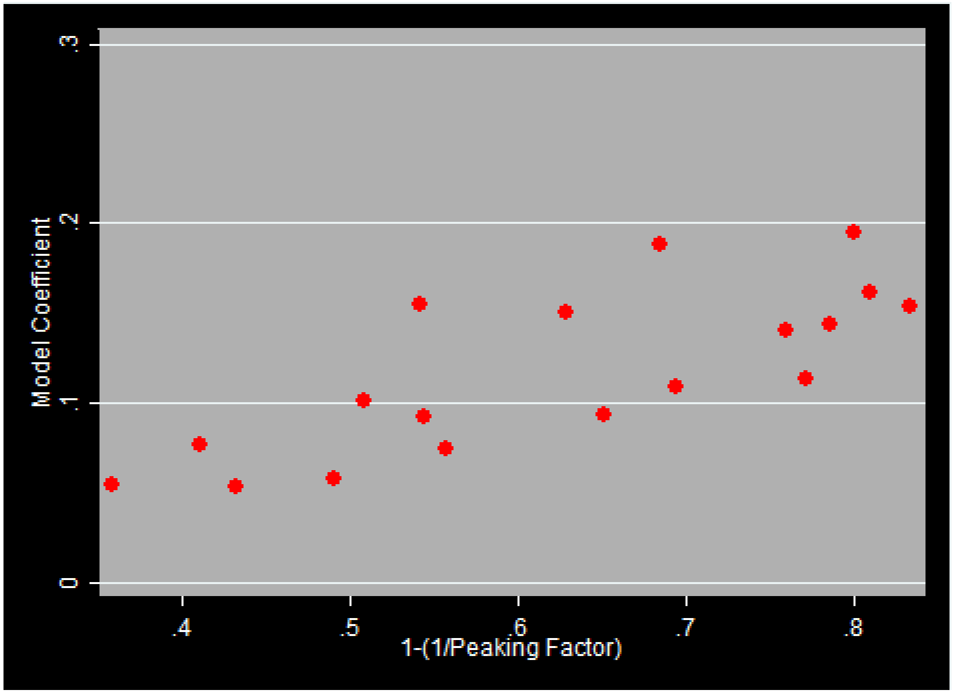

Figure 3 Supplier specific weather impacts as captured by rainfall adjusted reference ETo plotted by the transformed peaking factor

Table 4 Best fit linear regression of Figure 3's data

Dependent variable – Model coefficients in Figure 3

Independent Variable	**Coefficient (Standard Error)**	**t-statistic**
Transformed peaking factor	0.229 (0.048)	4.77†
Intercept	-0.026 (0.031)	-0.84‡

NOTE: Transformed peaking factor = 1-(1/peaking factor)
†Significant at 1% level
‡Insignificant

In the second set of supplier-specific models, we also examined how the impact of temperature and rainfall scales with respect to transformed peaking factors, since weather normalization on the basis of temperature and rainfall is also a question of interest. The plots (Figures 4 & 5) show that temperature effects scale linearly with respect to the transformed peaking factors, but rainfall has a weaker relationship albeit in the right general direction. This is a key reason why we think using rainfall adjusted reference ETo to weather normalize monthly production is a more reliable technique, since the effect of rainfall can be subsumed into the effect of reference ETo, both theoretically and empirically. But ultimately the difference between the two approaches is small, as sensitivity analyses presented later show.

The above discussion should also make clear why weather normalization on the basis of annual data is infeasible. Because peaking factors are a key input for our methodology, working with monthly data is necessary.

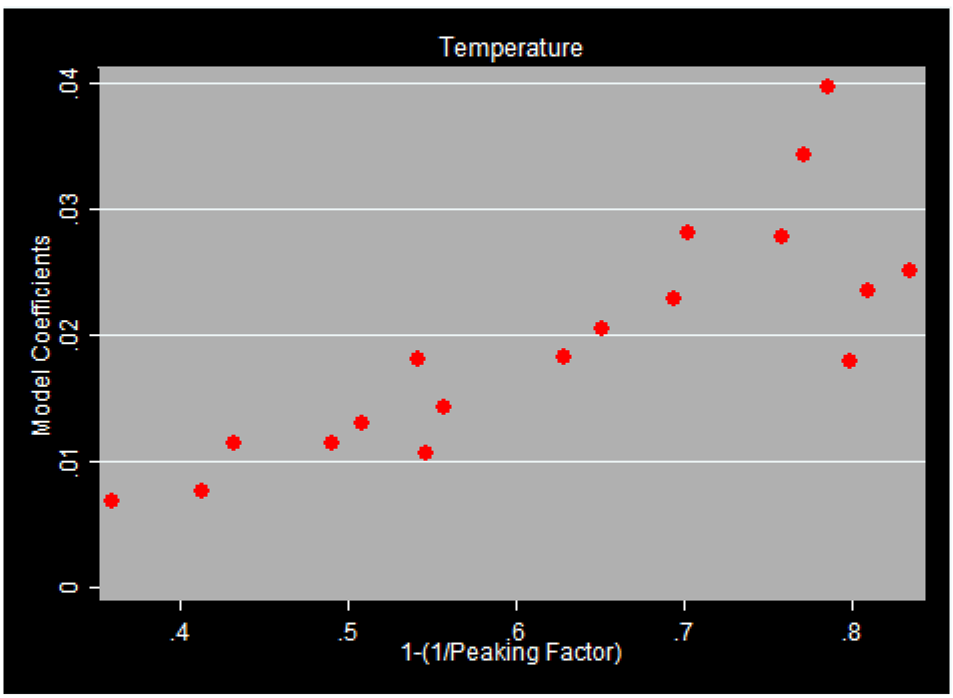

Figure 4 Supplier specific temperature impacts plotted by the transformed peaking factor

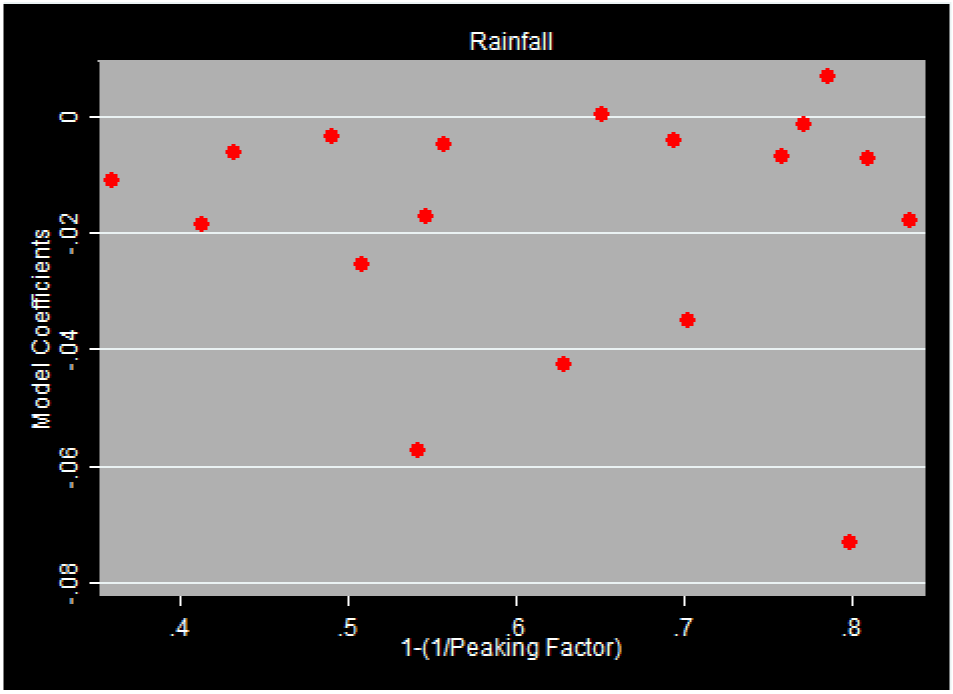

Figure 5 Supplier specific rainfall impacts plotted by the transformed peaking factors

4.2 The Estimated Impact of Weather

Having verified that the transformed peaking factor is a suitable scaling variable for adjusting weather response across different suppliers, we proceeded to estimate a single model that pools data across all 18 test suppliers (Appendix A describes the specification, estimation and sensitivity analyses of the pooled model in detail). A pooled model, because of greater sample size, allows for the estimation of weather impacts with greater precision. Furthermore, pooling of the data allows us to test for seasonal variation in weather response, which is difficult to do in a precise way using data only from a single supplier. These analyses suggest that a year can be broken into three seasonal groupings, November through March, April through June, and July through October. Weather impacts across months within a seasonal grouping were found to be comparable.

In the pooled models, the impact of weather was modeled in two different ways; one using temperature and rainfall; the other using rainfall adjusted reference ETo. All three weather measures were constructed using data obtained from PRISM.

As mentioned earlier, the independent effect of rainfall does not show as tight a relationship with transformed peaking factors, but temperature does. But, to weather normalize on the basis of temperature and rainfall requires using both variables in the model. In the reference ETo and rainfall model, we can choose to model these two effects independently, or combine them into a single measure—rainfall adjusted reference ETo. Given that the independent effect of rainfall does not scale tightly with transformed peaking factors, but the effect of rainfall adjusted reference ETo does, we feel the latter measure ought to be favored.

Table 5 shows results from the pooled statistical models. The table shows how deviations in temperature and rainfall, or rainfall adjusted reference ETo influence monthly production by season and peaking factor. So, for example, at a peaking factor of 2, a 1 degree positive deviation in temperature can be expected to raise monthly production by 1.26% during the months of November through March; by 1.32% during the months of April through June; and by 0.7% during the months of July through October. A greater than average rainfall has the opposite effect, lowering demand, hence production. A 1" greater rainfall in a month relative to what would be considered normal for that month reduces production by 0.31% if the month in question lies between November and March; by 3.86% during the months of April, May and June; and by 2.16% for the months July through October. Negative deviations would have the opposite effect. The magnitude of these weather impacts increases with peaking factor, but at a declining rate per our earlier discussion.

This pattern makes intuitive sense. Extra rain during the traditional rainy months has much less impact on perceptions than extra rain when one doesn't normally expect much. Similarly, unseasonably high temperatures have the greatest impact when one normally expects cool temperatures.

Table 5 also shows the impact of weather when weather is captured via rainfall adjusted reference ETo. These impacts are greater in percentage terms than the one for temperature because a 1"deviation in monthly reference ETo is of far greater consequence than a 1 degree deviation in monthly temperature. But, otherwise the pattern is comparable with weather deviations having the greatest impact during the spring season, just as they do in the temperature and rainfall scheme.

For deriving rainfall adjusted reference ETo, we subtracted 30% of monthly rainfall from monthly reference ETo, placing a floor of zero on the net result to prevent rainfall adjusted reference ETo from taking on negative values.[3] Assuming that 30% of monthly rainfall is effective gave us a better model fit in Figure 3 than the usual 20-25% assumption that landscape professionals use. This is not entirely surprising since we are talking about effective rainfall at the supplier level, not just in the context of irrigation. One can hypothesize that rainfall at the supplier level may

[3] Rainfall adjusted ET_o = maximum(0, (ET_o – 0.3 × rainfall))

have higher effectiveness since it probably also substitutes for sidewalk cleaning in addition to irrigation. Or, it may be that landscape professionals have traditionally underestimated the effect of rainfall.

Table 5 Estimated weather impacts by season and peaking factor

	Nov-Mar	Apr-Jun	Jul-Oct	Nov-Mar	Apr-Jun	Jul-Oct
Peaking factor	**Weather Normalization based upon temperature and rainfall**					
	Per 1° temperature deviation			**Per 1" rainfall deviation**		
1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
2	1.26%	1.32%	0.70%	-0.31%	-3.86%	-2.16%
3	1.69%	1.77%	0.93%	-0.42%	-5.11%	-2.87%
4	1.90%	1.99%	1.05%	-0.47%	-5.73%	-3.22%
5	2.03%	2.13%	1.12%	-0.50%	-6.10%	-3.44%
6	2.11%	2.22%	1.16%	-0.52%	-6.35%	-3.58%
	Weather normalization based upon rainfall adjusted reference ETo					
	Per 1" deviation					
1	0.00%	0.00%	0.00%			
2	5.26%	11.25%	6.03%			
3	7.07%	15.27%	8.13%			
4	7.99%	17.30%	9.19%			
5	8.54%	18.60%	9.83%			
6	8.91%	19.44%	10.26%			

4.3 Sensitivity Analyses Using Post-Baseline Years

In this section, we explore sensitivity of weather-normalized GPCD estimates with respect to alternative sources of weather data as well as alternative measures used to depict weather. This is done by predicting what production would have been during the three post-baseline years of 2006, 2007 and 2008 if weather during those years had been the same as it was during the baseline period. This requires working with Table 5's factors in reverse to correct actual production data. Appendix B describes these computations in greater detail.

We offer two comparisons. First we compare, using data from all 18 suppliers, how weather normalized GPCDs for the three post-baseline years compare when either temperature and rainfall, or rainfall adjusted reference ETo, is used to perform the weather normalization. In this first comparison, all weather data are drawn from PRISM, and the comparison tests for sensitivity to alternative measures used to depict weather. Then for the subset of 9 Phase I suppliers for whom we also have temperature and rainfall data from NOAA, we compare the performance of NOAA to PRISM data. Although it was shown earlier that temperature and rainfall data from NOAA and PRISM correlate strongly, this second sensitivity analysis offers further proof about the efficacy of using PRISM as a weather data source.

Table 6 shows the results of the first sensitivity analyses. In general, the results are comparable regardless of whether temperature and rainfall, or rainfall adjusted reference ETo, are used to perform the weather normalization. Sometimes one weather measure leads to a slightly higher estimate, and at other times, the other. There is no evidence of one weather measure consistently under or over-predicting relative to the other measure. That said, the normalization based upon rainfall adjusted reference ETo may be more reliable because the independent effects of rainfall do not scale as well as that of temperature (in the temperature and rainfall model) with respect to transformed peaking factors.

Table 7 shows the results of the second sensitivity analyses, the one that assesses the impact of using NOAA versus PRISM temperature and rainfall data. For generating this table the models were rerun on only the 9 Phase I suppliers, for each source of weather data. The model generated weather effects were then used to normalize production in the three post-baseline years. Except for Supplier #2, once again the weather normalized GPCDs are very close to one another, adding further weight to the acceptability of PRISM as a weather data source. For Supplier # 2, CIMIS data were used instead of NOAA data since that latter were unavailable, which may explain the slightly higher discrepancy.

Table 6 Sensitivity to alternate weather measures from PRISM

Supplier	Baseline GPCD	Peaking factor	Year	Actual GPCD	Temp. & rainfall normalized GPCD	Rainfall adjusted ref. ETo normalized GPCD
1			2006	195	199	201
1	209	2.9	2007	188	185	185
1			2008	195	189	189
2			2006	165	163	164
2	166	2.2	2007	173	170	171
2			2008	163	158	160
3			2006	121	124	124
3	124	1.8	2007	130	132	132
3			2008	128	129	129
4			2006	139	138	137
4	145	2.0	2007	141	142	141
4			2008	134	133	132
5			2006	142	140	140
5	154	1.7	2007	146	146	144
5			2008	140	138	138
6			2006	155	157	157
6	165	2.0	2007	152	152	150
6			2008	144	143	141
7			2006	175	170	167
7	181	2.2	2007	180	178	175
7			2008	166	161	159
9			2006	159	162	162
9	174	2.3	2007	164	164	163
9			2008	159	158	157
10			2006	122	121	121
10	133	1.6	2007	122	122	121
10			2008	110	109	109
11			2006	334	327	328
11	349	5.2	2007	338	329	327
11			2008	343	320	321
12			2006	313	317	318
12	323	4.4	2007	332	333	333
12			2008	342	336	334
13			2006	406	409	409
13	380	3.2	2007	480	479	476
13			2008	462	454	455

Supplier	Baseline GPCD	Peaking factor	Year	Actual GPCD	Temp. & rainfall normalized GPCD	Rainfall adjusted ref. ETo normalized GPCD
14			2006	210	211	212
14	266	6.0	2007	230	228	231
14			2008	240	234	232
15			2006	285	287	289
15	297	4.1	2007	270	267	270
15			2008	275	269	268
16			2006	444	445	448
16	516	4.7	2007	487	487	491
16			2008	498	487	484
18			2006	225	225	225
18	233	2.7	2007	229	232	231
18			2008	207	204	207
19			2006	383	384	381
19	411	5.0	2007	408	403	402
19			2008	408	405	407
20			2006	314	316	314
20	302	3.3	2007	321	317	315
20			2008	304	299	295

Table 7 Sensitivity to alternate sources of weather data (PRISM vs. NOAA)

Supplier	Baseline GPCD	Peaking factor	Year	Actual GPCD	Temp. & rainfall normalized GPCD (PRISM)	Temp. & rainfall normalized GPCD (NOAA)
1			2006	195	198	199
1	209	2.9	2007	188	185	185
1			2008	195	189	188
2			2006	165	162	164
2	166	2.2	2007	173	170	175
2			2008	163	159	160
3			2006	121	123	121
3	124	1.8	2007	130	132	130
3			2008	128	129	127
4			2006	139	138	138
4	145	2.0	2007	141	142	141
4			2008	134	133	132
5			2006	142	140	140
5	154	1.7	2007	146	146	146
5			2008	140	138	139
6			2006	155	157	157
6	165	2.0	2007	152	152	151
6			2008	144	143	142
7			2006	175	171	173
7	181	2.2	2007	180	178	178
7			2008	166	161	164
9			2006	159	161	161
9	174	2.3	2007	164	164	164
9			2008	159	159	157
10			2006	122	121	121
10	133	1.6	2007	122	122	122
10			2008	110	109	109

5. Conclusions

The analyses presented here test several questions pertaining to weather normalization. The first and foremost question was identifying a source of reliable weather data. We compare two different sources of weather data for this study, NOAA and PRISM. PRISM is a powerful new tool developed by the DWR that can provide temperature, rainfall, and reference ETo data from 1990 until present for all regions of California. Data are available for each 4x4 kilometer tile in California; there are over 26,000 such tiles covering the length and breadth of the state. There was a great deal of interest in assessing whether PRISM can serve as a one-stop shop for weather data. Temperature and rainfall data from these two sources show a very high level of correlation, suggesting that PRISM is indeed a reliable source of weather information.

Weather normalization is likely to proceed better if the same source of weather data is used for estimating the statistical models as for completing the normalization analyses for each supplier. Our analyses suggest that PRISM can serve as just such a source.

Second, there was a high degree of interest in assessing whether annual production can be reliably weather normalized since it would be an easy method to implement. But the statistical models failed to support this line of thinking since no statistically significant relationship could be detected between production and temperature or production and reference ETo at the annual level. The models, however, do offer strong support for attempting weather normalization at the monthly level. The approach described here shows the feasibility of using summer-to-winter peaking factors to capture the variation in weather's impact across different suppliers, obviating the need to collect detailed supplier characteristics.

We also compared two different approaches to weather normalization; one relying on temperature and rainfall; the other relying on rainfall adjusted reference ETo. Although the sensitivity analyses do not indicate large differences in the outcomes from either approach, we feel the latter approach is likely to prove more reliable in practice. The impact of rainfall adjusted reference ETo scales well with supplier peaking factors. In the temperature and rainfall approach, only temperature effects scale well with peaking factors, rainfall less so. Thus a methodology that does not require estimating the independent effect of rainfall on production is likely to be more reliable, and the weight of the empirical evidence presented earlier certainly supports this assertion.

The model estimated correction factors are not terribly sensitive to a few instances of bad weather data since they are estimated from multiple agencies. But, at the time of applying this methodology, one is working only with 12 data points per year. Bad data for one or two months could have a much greater influence on a specific agency's GPCD estimate, so great care needs to be exercised in this regard.

Appendix A Model specification and estimation

A.1 Model Specification

Household billing data are more complicated to model because of staggered billing cycles. These issues do not arise with monthly production data as calendar reporting periods can be made consistent across agencies. A key driver of production growth over time is population growth. At a first order of approximation, production growth over time ought to be proportional to population growth, suggesting a logarithmic relationship between the two variables. The basic monthly statistical model implied by such considerations is expressed in Eq. 1 below. The annual model, a simplified version of the monthly model, is discussed later.

$$Ln(prod)_{it} = \alpha + \beta Ln(pop)_{it} + \gamma_t(m_t) + \delta_t(w_t) + \theta_i + \varepsilon_{it} \quad \ldots. (1)$$

Where,

prod	stands for production for agency (i) in month (t)
pop	indicates corresponding population
m_t	a set of 12 indicators, one for each month
w_t	deviation in monthly weather relative to baseline average

And the model parameters have the following interpretation,

α	scaling constant, captures overall average
β	indicates percent increase in production for every 1% increase in population
γ_t	captures month-to-month variation in monthly production, assuming average weather prevails during every month
δ_t	captures variation in monthly production due to weather deviation from baseline averages
θ_i	indicator variable for agency (i) that captures variation in production levels across agencies
ε_{it}	model error

The primary parameters of interest are the coefficients on the weather deviation variables (δ_t), which yield the correction factors to be used for weather normalization.

The annual model is simply the aggregated version of the monthly model described above. In this case, annual production replaces monthly production on the left-hand-side of the equation, annual mean population replaces the monthly mean population, the 12 monthly indicators (m_t) are eliminated altogether, and the weather-deviation variables are derived at the annual instead of monthly level. So, instead of potentially having 12 monthly correction factors for temperature and 12 for rainfall, the annual model boils them down to 1 each. The annual model is much simpler in structure, thus easier to implement in practice—the question is does it work as well as the more detailed monthly model?

A.2 Estimation and Sensitivity Analyses

The basic structure of Eq. 1 can be refined and modified in several ways—and these refinements were tested during the sensitivity analysis phase to assess the robustness of the basic specification. Examples of the refinements that were tested include:

- Allowing the relationship between production and population (β) to vary by supplier instead of imposing a common relationship, and including additional time-trend variables for each supplier.
- Allowing the monthly indicator-variable coefficients (γ_t) to take different values by supplier instead of imposing a uniform seasonal pattern across all suppliers.
- Capturing weather deviations in two alternative ways; (1) using deviations in temperature and rainfall; and (2) using deviations in rainfall adjusted reference ETo.
- Testing whether logarithmically transforming weather variables before deriving the deviations from baseline averages improves the model fit.
- Assessing whether the correction factors (δ_t) vary significantly by month, in which case a total of 24 factors are required (12 for temperature and 12 for rainfall, or 12 for rainfall adjusted reference ETo), or whether these can be pared down to monthly groupings to simplify the methodology.
- Assessing whether weather impacts scale across different suppliers in a predictable way corresponding to a supplier's peaking factor.
- Testing for heteroscedasticity and autocorrelation in the model error (ε_{it}) and correcting for it.

A.3 Scaling Weather Impacts According to Peaking Factor

We recognize that each supplier's mix of weather-sensitive and weather-insensitive end uses is different; therefore, the impact of weather on total production cannot be identical across suppliers. Suppliers, however, do not always have good data to isolate these two types of end uses, so a weather normalization scheme that relies on the availability of such disaggregate data cannot succeed. Alternatively, one could collect supplier characteristics that correlate with weather-sensitive end-uses, such as irrigated landscape per capita, commercial air-conditioning, and so on, but these too are difficult to obtain in practice. Thus, we focused on a supplier's peaking factor as a way of scoring how suppliers rank relative to one another in terms of the proportion of total use that is accounted for by weather-sensitive end uses.

For peaking factor to work as a scoring variable, however, it must have a consistent relationship with the variation in weather impacts across suppliers. Whether such a relationship exists was tested in several ways by first running separate models for each supplier. These separate models were also used to estimate the peaking factors in a normal weather year. The model generated peaking factors correlate very highly with simply taking the ratio of highest and lowest production months per year, and then averaging these ratios across all years in the baseline. But the model generated peaking factors are conceptually cleaner, and improve the fit somewhat in Figure 3. The results of these analyses were discussed in Section 4. All the disparate pieces of evidence generated from these analyses bolster the case for using rainfall adjusted reference ETo to perform the weather normalization.

For deriving rainfall adjusted reference ETo, we subtracted 30% of rainfall from reference ETo, placing a floor of zero on the net result to prevent rainfall adjusted reference ETo from taking on negative values. Assuming 30% effective rainfall improved the model fit in Figure 3. While this effective rainfall parameter may appear somewhat higher than what most landscape professionals use, it must be remembered that we are talking about effective rainfall at the level of a supplier, not just in the context of irrigation. One can hypothesize that rainfall at the supplier level may have higher effectiveness since it probably also substitutes for sidewalk cleaning in addition to irrigation.

A.4 Model Results

Table 8 shows the key parameters of interest for the basic monthly model after correcting model error for autocorrelation and heteroscedasticty. This model is based on a total of 18 suppliers, 9 from phase I and 9 from phase II.

The fit of this model appears very good (adjusted R-square=0.95) and all the coefficients are statistically significant. The coefficient on the population variable is very close to 1, as we would expect: As per the estimated coefficient, 1% population growth led to 0.98% growth in production across these test agencies during the baseline period. The monthly indicator variables exhibit the expected pattern, with minimum production occurring in February and maximum in July and August. To interpret the monthly indicator coefficients in percentage terms they first need to be exponentiated. So, for example, on average February's production was 10.1% ($e^{-0.107}$-1) below January's (the reference month), while July's was 95.6% ($e^{0.671}$-1) above.

Next are shown the weather impact coefficients. The estimated coefficients on the weather deviation variables have been grouped into three seasonal categories because the magnitude of the monthly effects appeared similar within each category. The weather variables enter the models as a product of the weather variable and the transformed peaking factor to permit weather effects to vary across different suppliers. These coefficients also need to be exponentiated to give them a percentage interpretation, but for small coefficients this makes almost no difference. So, for example at a transformed peaking factor of 1, during the months of November through March, a 1 inch rise in rainfall adjusted reference ETo relative to what is considered normal leads to a 9.7% ($e^{0.093}$-1) increase in monthly production, and vice versa. During April through June, this relationship strengthens to 19.8% for every 1 inch deviation, while for the remainder of the year it weakens to roughly 12.1%. At other values of the transformed peaking factors, these would have to be multiplied by the coefficients before exponentiating.

Table 8 Estimated basic monthly model

Dependent variable – Ln(monthly production)

Variable	Coefficient	Std. error	t-statistic
Ln(population)	0.983	0.086	11.5
January	---		
February	-0.107	0.009	-11.3
March	0.082	0.014	6.0
April	0.230	0.014	17.0
May	0.434	0.014	30.4
June	0.554	0.014	40.5
July	0.671	0.012	54.3
August	0.665	0.012	54.6
September	0.559	0.012	46.6
October	0.437	0.012	37.0
November	0.156	0.011	13.6
December	0.047	0.010	4.9
Rainfall adjusted ref. ETo deviation x TPF‡ (Nov. through Mar.)	0.093	0.009	9.7
Rainfall adjusted ref. ETo deviation x TPF (Apr. through Jun.)	0.181	0.014	13.4
Rainfall adjusted ref. ETo deviation x TPF (Jul. through Oct.)	0.114	0.009	12.4
Constant	-5.15	0.868	-5.9
Adjusted R-square	0.95		

‡TPF, or transformed peaking factor, equals (1-(1/peaking factor))
NOTE: All coefficients are statistically significant at the 1% level. Agency specific fixed effects are included in the model. Error autocorrelation and heteroscedasticity corrected.

The coefficients reported in Table 8 are not the ones that we ultimately use for deriving weather-normalized GPCDs. Table 8 has been provided only to aid the reader's understanding of the model structure and interpretation of the coefficients. As mentioned earlier, the basic model was further relaxed by allowing the relationship between production and population, and the monthly pattern to vary by agency. This does not alter the weather coefficients much, but we should prefer these for the purpose of weather normalization. We also examined whether logarithmically transforming the weather variables improves the fit, but it did not do so, so we have not used that option as it is more difficult to explain and understand.

Table 9 Rainfall adjusted reference ETo coefficients from relaxed version of basic model using PRISM data

Dependent variable – Ln(monthly production)

Variable	Coefficient	Std. error	t-statistic
Rainfall adjusted ref. ETo deviation x TPF‡ (Nov. through Mar.)	0.102	0.006	16.1
Rainfall adjusted ref. ETo deviation x TPF (Apr. through Jun.)	0.213	0.009	24.4
Rainfall adjusted ref. ETo deviation x TPF (Jul. through Oct.)	0.117	0.006	18.6
Adjusted R-square	0.98		

‡TPF, or transformed peaking factor, equals (1-(1/peaking factor))
NOTE: All coefficients are statistically significant at the 1% level. Agency specific fixed effects are included in the model. Error autocorrelation and heteroscedasticity corrected.

Table 10 Temperature and rainfall coefficients from relaxed version of basic model using PRISM weather data

Dependent variable – Ln(monthly production)

Variable	Coefficient	Std. error	t-statistic
Temperature deviation x TPF‡ (Nov. through Mar.)	0.025	0.001	16.2
Temperature deviation x TPF (Apr. through Jun.)	0.026	0.002	15.6
Temperature deviation x TPF (Jul. through Oct.)	0.014	0.001	12.3
Rainfall deviation x TPF (Nov. through Mar.)	-0.006	0.002	-3.8
Rainfall deviation x TPF (Apr. through Jun.)	-0.079	0.006	-12.1
Rainfall deviation x TPF (Jul. through Oct.)	-0.044	0.004	-10.5
Adjusted R-square	0.98		

‡TPF, or transformed peaking factor, equals (1-(1/peaking factor))
NOTE: All coefficients are statistically significant at the 1% level. Agency specific fixed effects are included in the model. Error autocorrelation and heteroscedasticity corrected.

For the sake of brevity, Tables 9 and 10 only report weather coefficients from these relaxed versions of the basic model, which are then used in Section 4 to weather-normalize production data. Table 9 uses rainfall adjusted reference ETo to depict weather, while Table 10 uses temperature and rainfall. In both cases, weather data are taken from PRISM.

The coefficients in Table 9 are larger than those for temperature in Table 10 because a 1 inch deviation in monthly reference ETo is of far greater consequence than a 1 degree deviation in monthly temperature.

Figures 6 through 8 (based on Table 9's underlying model) compare model predictions to actual monthly production for a low, medium, and high peaking factor supplier. The high precision with which weather effects in Table 9 are estimated suggest that model predictions should compare well, and that is indeed borne out by these plots.

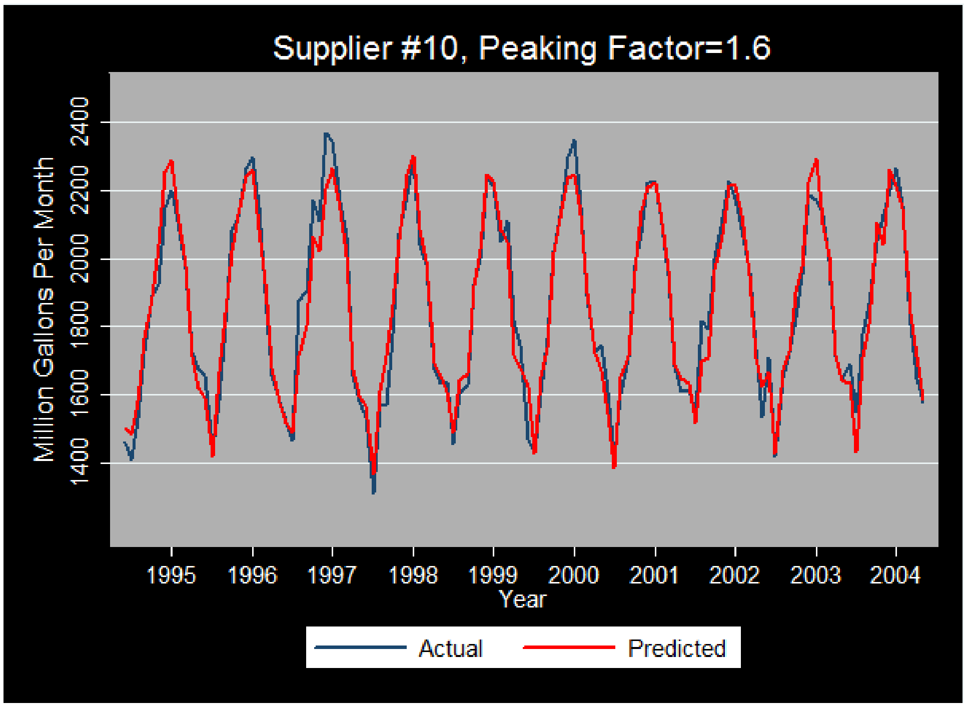

Figure 6 Actual vs. model prediction for a low peaking factor supplier

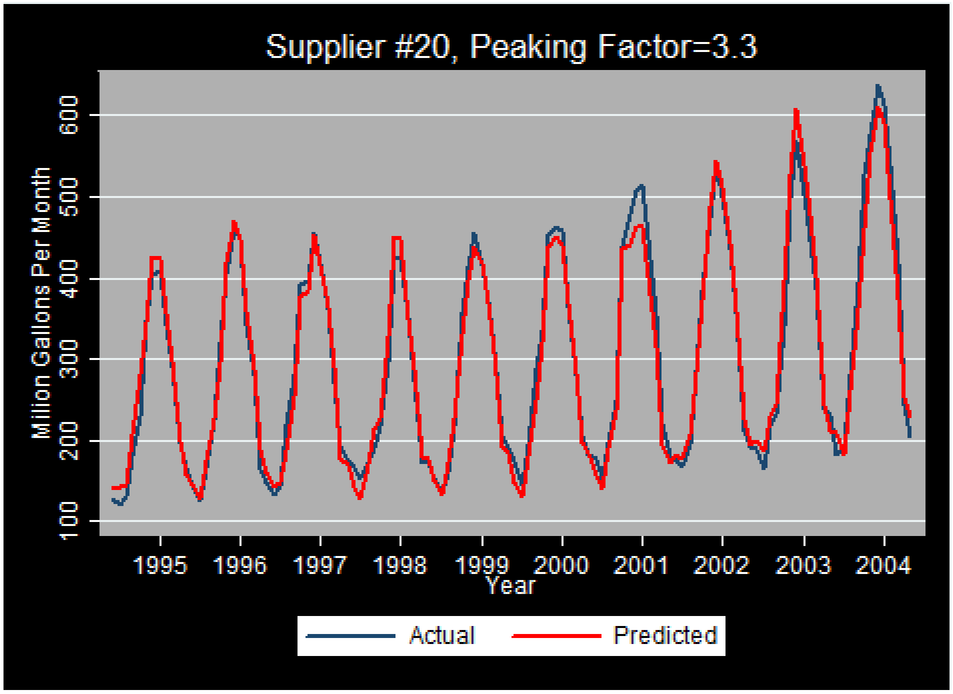

Figure 7 Actual vs. model prediction for a medium peaking factor supplier

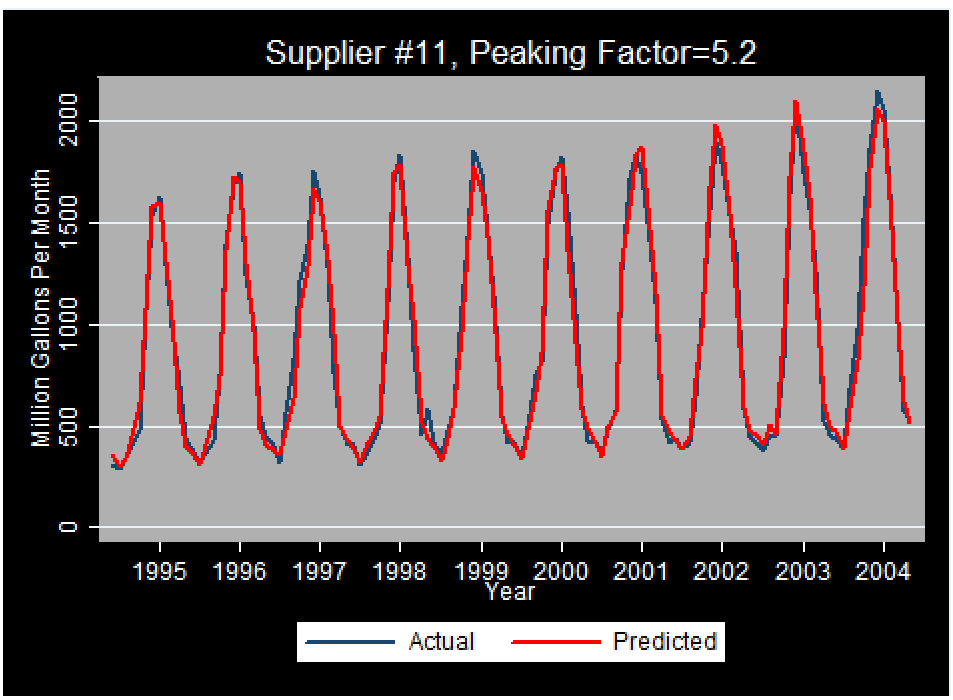

Figure 8 Actual vs. model prediction for a high peaking factor supplier

A.5 How Does the Annual Model Perform?

Given the salience of a supplier's peaking factor in determining its weather response, it is easy to see why an annual approach is infeasible. And without disaggregate data by month, or finer, it is impossible to estimate this factor. But, just to satisfy our curiosity we tried an annual model without including the peaking factor variable. This type of annual model was not successful in detecting a statistically significant relationship between production and temperature or production and reference ETo, although the effects of rainfall were significant. Therefore, weather normalization of annual data, while attractive for its simplicity, is not feasible.

Appendix B Applying the methodology: An example

To describe the computations involved, we have taken one year's data (2006) from one of the test agencies (Supplier #2) to illustrate the computations. We use the rainfall adjusted reference ETo approach for the purpose of this illustration. Table 11 shows rainfall adjusted reference ETo estimates by month for the year 2006, as well as averages from the baseline period, from which we derive the deviations between the compliance year and the baseline period.

In January of 2006, for example, rainfall adjusted reference ETo was 0.365 inches greater than the baseline average for this month. This positive weather deviation would have raised production in January of 2006, but by how much? That is obtained by taking the product of three variables, including the model coefficient for January, the transformed peaking factor, and the deviation between actual and baseline average weather as measured by rainfall adjusted reference ETo. This product is then exponentiated to express the impact of the weather deviation in terms of a multiplier since the models use the logarithmic transform of monthly production as the dependent variable. By what percentage did January's weather deviation increase production? This works out to 2.1% or a multiplier equal to 1.021 ($e^{0.543 \times 0.102 \times 0.365}$).

Had these deviations not occurred, our best estimate of what January 2006 production *would have been* is actual production scaled back by 2.1%. Thus, normalized production is obtained by dividing actual production by the impact multiplier.

Aggregating the weather-normalized monthly production estimates to the annual level and dividing by population and the number of days in the year then yields the weather-normalized GPCD.

Table 11 Monthly weather normalization: An example of the computations involved

Month	Actual rainfall adjusted reference ETo	Baseline average rainfall adjusted reference ETo	Deviation in rainfall adjusted reference ETO	Peaking factor	Transformed peaking factor	Model Coefficient	Weather deviation impact multiplier	Actual production	Normalized production
1	1.332	0.967	0.365	2.19	0.543	0.102	1.021	1170	1146
2	1.202	0.937	0.265	2.19	0.543	0.102	1.015	1200	1183
3	1.186	2.606	-1.419	2.19	0.543	0.102	0.924	1038	1123
4	2.398	3.790	-1.392	2.19	0.543	0.213	0.851	1063	1249
5	4.644	5.100	-0.456	2.19	0.543	0.213	0.949	1506	1588
6	6.609	5.907	0.703	2.19	0.543	0.213	1.085	1876	1730
7	7.665	6.875	0.790	2.19	0.543	0.117	1.052	2256	2145
8	6.577	6.376	0.201	2.19	0.543	0.117	1.013	2097	2070
9	5.163	4.848	0.315	2.19	0.543	0.117	1.020	1994	1954
10	3.510	3.174	0.336	2.19	0.543	0.117	1.022	1721	1685
11	2.523	1.850	0.673	2.19	0.543	0.102	1.038	1536	1480
12	1.651	1.055	0.596	2.19	0.543	0.102	1.034	1310	1268

NOTES: Transformed peaking factor = 1-(1/Peaking factor)
Model coefficients are taken from Table 9.